

A Multivariate Approach for Checking Resiliency in Access Control^{*}

Jason Crampton, Gregory Gutin, and Rémi Watrigant

Royal Holloway University of London

Abstract. In recent years, several combinatorial problems were introduced in the area of access control. Typically, such problems deal with an authorization policy, seen as a relation $UR \subseteq U \times R$, where $(u, r) \in UR$ means that user u is authorized to access resource r . Li, Tripunitara and Wang (2009) introduced the RESILIENCY CHECKING PROBLEM (RCP), in which we are given an authorization policy, a subset of resources $P \subseteq R$, as well as integers $s \geq 0$, $d \geq 1$ and $t \geq 1$. It asks whether upon removal of any set of at most s users, there still exist d pairwise disjoint sets of at most t users such that each set has collectively access to all resources in P . This problem possesses several parameters which appear to take small values in practice. We thus analyze the parameterized complexity of RCP with respect to these parameters, by considering all possible combinations of $|P|, s, d, t$. In all but one case, we are able to settle whether the problem is in FPT, XP, W[2]-hard, para-NP-hard or para-coNP-hard. We also consider the restricted case where $s = 0$ for which we determine the complexity for all possible combinations of the parameters.

1 Introduction

1.1 Context and definition of the problem

Access control is a fundamental aspect of the security of any multi-user computing system. Typically, it is based on the idea of specifying and enforcing an authorization policy, identifying which interactions between a set of users U and a set of resources R are to be allowed by the system [11]. More formally, an authorization policy is defined as a relation $UR \subseteq U \times R$, where $(u, r) \in UR$ means that user u is authorized to access resource r . Quite recently, we have seen the introduction of resiliency policies, whose satisfaction indicates that a system will continue to function as intended in the absence of some number of authorized users [10, 12]. Li, Tripunitara and Wang’s seminal work [10] introduces a number of problems associated with the satisfaction of a resiliency policy. One of their motivating examples concerns the emergency response to a natural disaster, where teams of users must perform the same critical operation(s) at multiple (distinct) geographical locations. Thus the members of each

^{*} This research was partially supported by EPSRC grant EP/K005162/1. Gutin’s research was also supported by Royal Society Wolfson Research Merit Award.

team must be authorized collectively to perform the operation(s). In addition, we may wish to impose an upper bound on the size of the teams because, for example, of constraints on transportation.

For a user $u \in U$ and a set of users $V \subseteq U$, we define $N_{UR}(u) = \{r \in R : (u, r) \in UR\}$ the *neighborhood* of u and, by extension, $N_{UR}(V) = \bigcup_{u \in V} N_{UR}(u)$ the *neighborhood* of V , omitting the subscript UR if the authorization policy is clear from the context. Given an authorization policy $UR \subseteq U \times R$, an instance of the RESILIENCY CHECKING PROBLEM (RCP) is defined by a resiliency policy $\text{res}(P, s, d, t)$, where $P \subseteq R$, $s \geq 0$, $d \geq 1$ and $t \geq 1$. We say that UR *satisfies* $\text{res}(P, s, d, t)$ if and only if for every subset $S \subseteq U$ of at most s users, there exist d pairwise disjoint subsets of users V_1, \dots, V_d such that for all $i \in \{1, \dots, d\}$:

$$V_i \cap S = \emptyset, \quad (1)$$

$$|V_i| \leq t, \quad (2)$$

$$N(V_i) \supseteq P. \quad (3)$$

We are now ready to define the main problem we study in this paper:

Resiliency Checking Problem (RCP)

Input: $UR \subseteq U \times R$, $P \subseteq R$, $s \geq 0$, $d \geq 1$, $t \geq 1$.

Question: Does UR satisfy $\text{res}(P, s, d, t)$?

Furthermore, we will adopt the bracket notation $\text{RCP}\langle \rangle$ used by Li *et al.* [10] to denote some restrictions of the problem, in which one or more parameters (among s , d and t) are fixed. In particular, we will consider the cases where s and d are respectively set to 0 and/or 1 (or other fixed positive values), while t might be set to ∞ , meaning that there is no constraint on the size of the sets (which is actually equivalent to $t = |P|$), implying that we may assume in the remainder that $t \leq |P|$). For instance, $\text{RCP}\langle s = 0 \rangle$ denotes the variant in which s is fixed to 0, *i.e.* we ask for the satisfaction of $\text{res}(P, 0, d, t)$. In the remainder of the paper, we set $p = |P|$.

Given an instance of $\text{RCP}\langle \rangle$, we say that a set of d pairwise disjoint subsets of users $V = \{V_1, \dots, V_d\}$ satisfying conditions (2) and (3) is a *set of teams*. For such a set of teams, we define $\mathcal{U}(V) = \bigcup_{i=1}^d V_i$. Given $U' \subseteq U$, the *restriction* of UR to U' is defined by $UR|_{U'} = UR \cap (U' \times R)$. Finally, a set of users $S \subseteq U$ is called a *blocker set* if for every set of teams $V = \{V_1, \dots, V_d\}$, we have $\mathcal{U}(V) \cap S \neq \emptyset$. Equivalently, observe that S is a blocker set if and only if $UR|_{U \setminus S}$ does not satisfy $\text{res}(P, 0, d, t)$. Throughout the paper, we write $[d]$ to denote $\{1, \dots, d\}$ for any integer $d \geq 1$, and we will often make use of the $O^*(\cdot)$ notation, which omits polynomial factors and terms.

1.2 Parameters

An instance of $\text{RCP}\langle \rangle$ contains several parameters (namely s , d and t) which may be used for the complexity analysis of the problem. An interesting point of the work of Li *et al.* [10] is that the number of users in an organization will

typically be large in comparison to the other parameters (s , d , t , and even p) in practice. In their experiments, the maximum values used are $n = 100$, $p = 10$ and $d = 7$ (they only run experiments on the variant where $t = \infty$, but, as we observed previously, we may set $t = p$). With this in mind, we exploit the theory of fixed-parameter tractability in order to settle the parameterized complexity of the problem.

Given an instance x (of size $|x|$) of a decision problem, with some parameter¹ k , we are interested in algorithms deciding whether x is positive or negative in polynomial time when k is bounded above by a constant. More precisely, if such an algorithm has running time $O(f(k)|x|^{O(1)})$ for some computable function f , then we will say that this algorithm is fixed-parameter tractable (FPT), while if its running time is $O(|x|^{f(k)})$ for some computable function f , we will say that this algorithm is XP (an FPT algorithm is thus an XP algorithm). By extension, FPT (resp. XP) gathers all problems for which an FPT (resp. XP) algorithm exists. Proving the NP-hardness of a problem in the case where a parameter k is bounded above by a constant immediately forbids the existence of any XP (and thus FPT) algorithm unless $P = NP$. In this case, we will say that this parameterized problem is para-NP-hard. A similar definition can be given using coNP-hard and coNP instead of NP-hard and NP, respectively, leading to the para-coNP-hard complexity class (and thus, if a problem is shown to be para-coNP-hard, then it does not belong to XP unless $P = \text{coNP}$). In the following, para-(co)NP-hard denotes the union of para-NP-hard and para-coNP-hard. Finally, the $W[i]$ -hierarchy of parameterized problems is typically used to rule out the existence of an FPT algorithm, under the widely believed conjecture that $\text{FPT} \neq W[1]$. For more details about fixed-parameter tractability, we refer the reader to the recent monographs [2, 4].

1.3 Related work

As one might expect, the $\text{RCP}\langle \rangle$ problem is strongly related to some known combinatorial problems. Indeed, one can observe that $\text{RCP}\langle s = 0, d = 1 \rangle$ is equivalent to the SET COVER problem, while $\text{RCP}\langle s = 0, t = \infty \rangle$ can be reduced in a straightforward way from the DOMATIC PARTITION problem (in the DOMATIC PARTITION problem, one asks whether a given graph admits k pairwise disjoint dominating sets). Li *et al.* [10] obtained several (mainly negative) results for $\text{RCP}\langle \rangle$ in some restricted cases which can be summarized by the following theorem.

Theorem 1 ([10]). *We have the following:*

- $\text{RCP}\langle \rangle$, $\text{RCP}\langle d = 1 \rangle$ and $\text{RCP}\langle t = \infty \rangle$ are NP-hard and are in² coNP^{NP} ;
- $\text{RCP}\langle s = 0, d = 1 \rangle$, $\text{RCP}\langle s = 0, t = \infty \rangle$ are NP-hard;

¹ Note that one can aggregate several parameters p_1, \dots, p_m by defining $k = p_1 + \dots + p_m$, in which case we will say the parameter is (p_1, \dots, p_m) .

² coNP^{NP} is the set of problems whose complement can be solved by a non-deterministic Turing machine having access to an oracle to a problem in NP.

– $\text{RCP}\langle d = 1, t = \infty \rangle$ can be solved in linear time.

In addition, they developed and implemented an algorithm for $\text{RCP}\langle \rangle$ which consists of (i) enumerating all subsets of at most s users, and (ii) for each such subset S , determining the satisfaction of $\text{res}(P, 0, d, t)$ for $UR|_{U \setminus S}$. Step (ii) is achieved by a SAT formulation of the problem and the use of an off-the-shelf SAT solver, while they develop a pruning strategy in order to avoid the entire enumeration of all subsets of users of size at most s , resulting in an efficient speed-up of step (i). Quite surprisingly, they observe that the bottleneck of their algorithm lies in the second step, where an instance of $\text{RCP}\langle s = 0 \rangle$ has to be solved. This motivated us to focus on the parameterized complexity of $\text{RCP}\langle s = 0 \rangle$ separately.

1.4 Contribution and organization of the paper

Our goal in this paper is thus to determine the parameterized complexity of $\text{RCP}\langle \rangle$ and $\text{RCP}\langle s = 0 \rangle$ with respect to parameters p, s, d, t , by considering every possible combination of them. In each case, we aim at determining whether the problem is (i) in FPT, (ii) in XP but $W[i]$ -hard for some $i \geq 1$, or (iii) para-(co)NP-hard.

Figure 1 summarizes the (already known and) obtained results for $\text{RCP}\langle \rangle$ and $\text{RCP}\langle s = 0 \rangle$ with respect to all possible combinations of the parameters

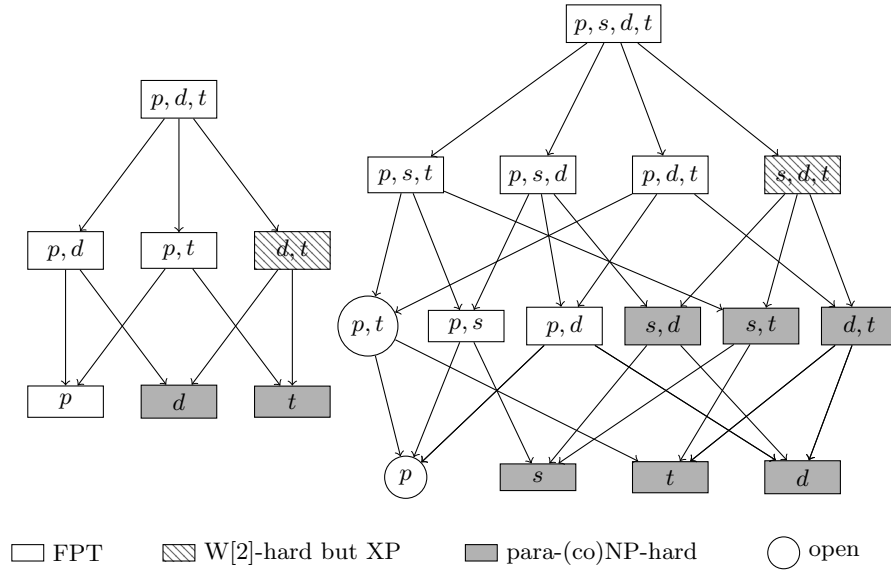


Fig. 1: Schemas of the complexity of $\text{RCP}\langle s = 0 \rangle$ (left) and $\text{RCP}\langle \rangle$ (right) after the results obtained in this paper (see the end of this section for the difference between old and new results).

specified previously. An arrow $A \longrightarrow B$ means that A is a larger parameter than B , in the sense that an FPT algorithm parameterized by B implies an FPT algorithm parameterized by A , and, conversely, any negative result parameterized by A implies the same negative result parameterized by B . Since (under classical complexity assumptions) a decision problem is either in one of the previous cases (i), (ii) or (iii), one can observe that the parameterized complexity of $\text{RCP}\langle s = 0 \rangle$ is now completely determined with respect to all possible combinations of parameters p , d and t . Concerning the more general case $\text{RCP}\langle \rangle$, only the parameterization by p only remains unknown (recall that as we mentioned earlier, we may assume in any instance that $t \leq p$, implying that adding t in the parameter list is of no importance concerning the membership in these complexity classes, both for positive or negative results).

The next section gathers all our results for the general case $\text{RCP}\langle \rangle$, namely:

- membership in XP parameterized by (s, d, t) (Theorem 2),
- membership in FPT parameterized by (p, d) or (p, s) (Theorem 3),
- para-coNP-hardness parameterized by (d, t) (Theorem 4),
- para-NP-hardness parameterized by (s, t) (Theorem 5).

Note that the para-NP-hardness for (s, d) was already known (Theorem 1), as well as the $W[2]$ -hardness for (s, d, t) (see explanation in Section 2.2).

Section 3 gathers all our results for the restricted case $\text{RCP}\langle s = 0 \rangle$, namely:

- an FPT algorithm parameterized by (d, p) with an optimal running time (under ETH) when d is fixed (Theorems 6 and 7),
- membership in FPT parameterized by p only (Theorem 8).

Note that the $W[2]$ -hardness for (d, t) is inherited from $\text{RCP}\langle \rangle$, while the XP membership results from a brute-force enumeration of all subsets of users of size dt . We also investigate in this section the question of data (user) reductions and present positive and negative kernelization results depending on the considered variant: $\text{RCP}\langle s = 0 \rangle$ or $\text{RCP}\langle s = 0, t = \infty \rangle$ (Theorem 10). We finally conclude the paper in Section 4.

2 The general case

2.1 Positive results

First, observe that there exists a simple XP algorithm for $\text{RCP}\langle \rangle$ parameterized by (s, d, t) . Indeed, recall that the problem actually aims to check whether there is a set $S \subseteq U$ of size at most s such that for any set of teams $V = \{V_1, \dots, V_d\}$ we have $S \cap \mathcal{U}(V) \neq \emptyset$, and note that finding a set of teams is exactly the $\text{RCP}\langle s = 0 \rangle$ problem, which is in XP parameterized by (d, t) , as said in Section 1.4. Hence, since $|\mathcal{U}(V)| \leq dt$, by finding iteratively a set of teams and branching on each element to be removed from it (and included in the future blocker set), one can determine whether there exists a blocker set of size at most s in XP time parameterized by (s, d, t) :

Theorem 2. *$\text{RCP}\langle \rangle$ is in XP when parameterized by (s, d, t)*

Despite its simplicity, this result is actually somehow tight. First, as we will see later (Section 2.2), $\text{RCP}\langle \rangle$ is $\text{W}[2]$ -hard with this parameterization. In addition, considering a strict subset of $\{s, d, t\}$ as a parameter makes the problem para-(co)NP-hard (Theorems 1, 4 and 5). A way of going further is to “replace” t by p (since we may assume $t \leq p$). With this modification, we show in the next result how to get rid of the parameter s or d by designing an FPT algorithm parameterized by (p, d) or (p, s) .

Theorem 3. $\text{RCP}\langle \rangle$ is FPT when parameterized by $(p, \min\{s, d\})$.

Proof. Without loss of generality, we may assume $P = R$ as well as $N(u) \neq \emptyset$ for all $u \in U$. For all $C \subseteq P$, let $U_C = \{u \in U : N(u) = C\}$ (notice that we might have $U_C = \emptyset$ for some $C \subseteq P$). Let $S \subseteq U$ be a blocker set of size at most s , *i.e.* a set whose removal makes $\text{res}(P, 0, d, t)$ unsatisfiable. Moreover, assume that S is a *minimal blocker set*, meaning that there does not exist $S' \subsetneq S$ such that the removal of S' makes $\text{res}(P, 0, d, t)$ unsatisfiable.

Claim. For all $C \subseteq P$, $U_C \cap S \neq \emptyset$ implies that $|U_C \setminus S| < d$.

Before proving the claim, notice that for all $u \in U_C \cap S$, there exists a set of teams $V = \{V_1, \dots, V_d\}$ such that (i) $\mathcal{U}(V) \cap S = \{u\}$, and (ii) $|\mathcal{U}(V) \cap U_C| \leq d$. Condition (i) comes from the minimality of S , while Condition (ii) comes from the fact that otherwise, there would exist $i \in [d]$ such that $|V_i \cap U_C| \geq 2$, and removing one user from V_i , arbitrarily chosen in $(V_i \cap U_C) \setminus \{u\}$, produces another set of teams V' with $\mathcal{U}(V') \subsetneq \mathcal{U}(V)$ (with exactly one element less) and still such that $V \cap S = \{u\}$. Applying this strategy iteratively, we can get a set of teams V as desired.

Proof (of the claim). To do so, let $u \in U_C \cap S$ and $V = \{V_1, \dots, V_d\}$ defined as previously. If we have $|U_C \setminus S| \geq d$, then there exists $v \in U_C \setminus S$ such that $v \notin \mathcal{U}(V)$ (since $|\mathcal{U}(V) \cap U_C| \leq d$, and since $u \in S \cap U_C$, it follows that $|(U_C \setminus S) \cap \mathcal{U}(V)| \leq d - 1$), in which case we have that $(\mathcal{U}(V) \setminus \{u\}) \cup \{v\}$ is the union of a set of teams which does not intersect S (recall that $\mathcal{U}(V) \cap S = \{u\}$), and satisfies $\text{res}(P, 0, d, t)$ (since $N(u) = N(v)$), a contradiction. \square

We now define a reduced set of users $U^r \subseteq U$ composed of $d_C = \min\{|U_C|, d\}$ users from U_C chosen arbitrarily, for all $C \subseteq P$. By construction, observe that $|U^r| \leq d2^p$. We also define, for all $C \subseteq P$, $U_C^r = U_C \cap U^r$. Finally, consider an algorithm which outputs that $\text{res}(P, s, d, t)$ is unsatisfiable if and only if there exists a blocker set $S \subseteq U^r$ of the instance induced by U^r (*i.e.* with authorization policy $UR|_{U^r}$), and such that $\sum_{C \subseteq P} \zeta_S(C) \leq s$, where

$$\zeta_S(C) = \begin{cases} |S \cap U_C^r| + |U_C| - d_C & \text{if } S \cap U_C^r \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

in which case we will say that S is a *reduced blocker set*. We will prove that this algorithm is FPT parameterized by $(p, \min\{s, d\})$, and is correct.

Concerning the running time, observe first that the construction of U^r as well as the evaluation of ζ_S , given $S \subseteq U^r$, takes $O^*(2^p)$ time. Then, for any reduced blocker set $S \subseteq U^r$, notice that $|S \cap U_C^r| \leq \min\{s, d\}$ for all $C \subseteq P$, and that any set $S' \subseteq U^r$ such that $|S' \cap U_C^r| = |S \cap U_C^r|$ for all $C \subseteq P$ is also a reduced blocker set (because $N(u) = N(v)$ for all $u, v \in C$, for all $C \subseteq P$). Hence, instead of enumerating every possible subset S of U^r , it is sufficient to enumerate the sizes of each intersection with U_C^r for all $C \subseteq P$, and pick the right number of users in U_C^r in an arbitrary way. Since its intersection is of size at most $\min\{s, d\}$, the number of sets to enumerate is $O((\min\{s, d\} + 1)^{2^p})$. Then, for each obtained set $S \subseteq U^r$, we can check whether it is a blocker set of $UR|_{U^r}$ by solving the $\text{RCP}\langle s = 0 \rangle$ problem on the instance $UR|_{U^r \setminus S}$ in FPT time parameterized by p (using, *e.g.*, Theorem 8).

It now remains to prove its correctness, by proving that there exists a reduced blocker set if and only if $\text{res}(P, s, d, t)$ is unsatisfiable. If such a set S exists, then define, for each $C \subseteq P$, a set $S_C \subseteq U_C$ composed of $S \cap U_C^r$ plus all users in $U_C \setminus U_C^r$. By construction, $|S_C| = \zeta_S(C)$, and thus $S^* = \bigcup_{C \subseteq P} S_C$ contains at most s users. We now prove that S^* is a blocker set: suppose by contradiction that there exists a set of teams $V = \{V_1, \dots, V_d\}$ such that $\mathcal{U}(V) \cap S^* = \emptyset$. As we saw previously, we may assume that $|V_i \cap U_C| \leq 1$ for all $i \in [d]$ and all $C \subseteq P$. Let $I_V = \{i \in [d] : V_i \cap (U_C \setminus U_C^r) \neq \emptyset\}$. We show that we can turn V into another set of teams V' such that $\mathcal{U}(V') \subseteq U^r$ (*i.e.* such that $I_{V'} = \emptyset$), implying that S is not a reduced blocker set, a contradiction. If $I_V = \emptyset$, then we are done. Otherwise let $i \in I_V$ and $u \in V_i \cap (U_C \setminus U_C^r)$. By construction of U^r , there exists $v \in U_C^r$, and thus $(V \setminus \{u\}) \cup \{v\}$ is the union of a set of teams V' (recall that $N(u) = N(v)$) such that $i \notin I_{V'}$. Repeating this transformation at most d times, we naturally obtain a set of teams V' such that $I_{V'} = \emptyset$ as desired.

Conversely, suppose that $\text{res}(P, s, d, t)$ is unsatisfiable, *i.e.* there exists a blocker set of users $S \subseteq U$ of size at most s . As previously, we may assume that S is a minimal blocker set. We now use the previous Claim, and thus for all $C \subseteq P$, $|S \cap U_C| \geq \max\{0, |U_C| - d + 1\}$. Thus, we may assume, without loss of generality (since, again, $N(u) = N(v)$ for all $u, v \in U_C$) that $U_C \setminus U_C^r \subseteq S$. Then, we define $S^r = S \setminus (\bigcup_{C \in \wp(C)} U_C \setminus U_C^r)$. Observe that for all $C \subseteq P$, we have:

$$\begin{aligned} \zeta_{S^r}(C) &= |S^r \cap U_C^r| + |U_C| - d_C \\ &= |S^r \cap U_C^r| + |U_C \setminus U_C^r| \\ &= |S \cap U_C| \end{aligned}$$

and thus $\sum_{C \subseteq P} \zeta_{S^r}(C) = \sum_{C \subseteq P} |S \cap U_C| = |S| \leq s$. Finally, S^r is indeed a blocker set of the instance induced by U^r , since otherwise, there would exist a set of teams $V = \{V_1, \dots, V_d\}$ with $\mathcal{U}(V) \subseteq U^r$ such that $\mathcal{U}(V) \cap S^r = \emptyset$, which would imply that $\mathcal{U}(V) \cap S = \emptyset$ as well, a contradiction. \square

2.2 Negative results

It is worth pointing out that the reduction of [10, Lemma 3] proving the NP-hardness of $\text{RCP}\langle s = 0, d = 1 \rangle$ actually proves the W[2]-hardness of this problem

parameterized by t (from SET COVER parameterized by the size of the solution [4]). Another implication of this reduction is the para-NP-hardness of $\text{RCP}\langle \rangle$ when parameterized by (s, d) . We now complement this result by showing that $\text{RCP}\langle d = 1, t = \tau \rangle$ is coNP-hard for every fixed $\tau \geq 3$, implying para-coNP-hardness of $\text{RCP}\langle \rangle$ parameterized by (d, t) . The result is obtained by a reduction from the δ -HITTING SET problem for every $\delta \geq 2$.

Theorem 4. $\text{RCP}\langle d = 1, t = \tau \rangle$ is coNP-hard for every fixed $\tau \geq 3$.

Proof. We reduce from the δ -HITTING SET problem, in which we are given a ground set $V = \{v_1, \dots, v_n\}$, a set $S = \{S_1, \dots, S_m\}$ with $S_j \subseteq V$ and $|S_j| = \delta$ for all $j \in [m]$ and an integer k , and where the goal is to find a set $C \subseteq V$ of size at most k and such that $C \cap S_j \neq \emptyset$ for all $j \in [m]$. This problem is known to be NP-hard for every $\delta \geq 2$ [6].

Hence, let (V, S, k) be an instance of δ -HITTING SET defined as above. For every $j \in [m]$, fix an arbitrary ordering of S_j , which can thus be seen as a tuple $(v_{i_1}, \dots, v_{i_\delta})$, allowing us to define $S_j[x] = v_{i_x}$ for all $x \in [\delta]$.

We define a set of users $U = U^V \cup U^S$, where $U^V = \{u_1^V, \dots, u_n^V\}$ and $U^S = \{u_1^S, \dots, u_m^S\}$. We then define a set of resources $R = R^V \cup R^S \cup \{r^*\}$, where $R^S = \bigcup_{j=1}^m P^j$ with $P^j = \{p_1^j, \dots, p_\delta^j\}$ for all $j \in [m]$, and where R^V contains one resource r_Q^V for every subset Q of $\delta - 1$ users of U^V .

We now define the authorization policy UR by giving $N(u)$ for every $u \in U$. For every $i \in [n]$, $N(u_i^V)$ is composed of $\{p_x^j : j \in [m], x \in [\delta] \text{ such that } S_j[x] = v_i\}$ together with all resources r_Q^V such that $u_i^V \notin Q$, for every subset Q of $\delta - 1$ users of U^V . For all $j \in [m]$, $N(u_j^S)$ is composed of r^* together with $R^S \setminus P^j$. To conclude the construction, we let $P = R$, $t = \delta + 1$, $d = 1$, and $s = k$. Clearly this reduction can be done in polynomial time.

The remainder consists in proving that every team (*i.e.* sets of at most t users having collectively access to all R) is of the form $T_j = \{u_{i_1}^V, \dots, u_{i_\delta}^V, u_j^S\}$ such that $S_j = \{v_{i_1}, \dots, v_{i_\delta}\}$. If this is true, then observe that since, for every $j \in [m]$, user u_j^S only belongs to team T_j , we will be able to suppose *w.l.o.g.* that it does not belong to any blocker set, and thus the set of all teams will be in one-to-one correspondance with the sets in S , implying that the obtained instance has a blocker set of size at most s ($= k$) if and only if there is a hitting set of size at most k .

Let $T \subseteq U$ of size at most t . By construction, we need at least δ users from U^V to have access to all resources in R^V (indeed, every set Q of $\delta - 1$ users from U^V has only access to $R^V \setminus \{r_Q^V\}$), and we also need at least one user from U^S to have access to r^* . Hence, $|T \cap U^V| = \delta$ and $T \cap U^S = \{u_j^S\}$ for some $j \in [m]$. Now, notice that u_j^S has access to all resources in R but P^j , which implies that $T \cap U^V$ must have collectively access to all resources in P^j . However, this can only happen if $T \cap U^V = \{u_{i_1}^V, \dots, u_{i_\delta}^V\}$, where $S_j = \{v_{i_1}, \dots, v_{i_\delta}\}$, concluding the proof. \square

We also settle the case of $\text{RCP}\langle \rangle$ parameterized by (s, t) (and thus $\text{RCP}\langle s = 0 \rangle$ parameterized by t). The result is obtained by a reduction from the 3-DIMENSIONAL MATCHING problem.

Theorem 5. $\text{RCP}\langle s = 0, t = 4 \rangle$ is NP-hard.

Proof (of Theorem 5). We reduce from the 3-DIMENSIONAL MATCHING problem, in which we are given three sets X , Y and Z of n elements each, a set $M \subseteq X \times Y \times Z$ of hyperedges, and an integer k . The goal is to find $M' \subseteq M$ with $|M'| \geq k$ such that $\forall e, e' \in M'$ with $e \neq e'$, $e = (x, y, z)$, $e' = (x', y', z')$, we have $x \neq x'$, $y \neq y'$ and $z \neq z'$ (in that case, we will say that these two hyperedges are *disjoint*). We note $m = |M|$, $X = \{x_1, \dots, x_n\}$, $Y = \{y_1, \dots, y_n\}$, $Z = \{z_1, \dots, z_n\}$, and $M = \{e_1, \dots, e_m\}$.

We then define the following set of resources:

$$\begin{aligned} P = & \{r_1^X, \dots, r_m^X\} \\ & \cup \{r_1^Y, \dots, r_m^Y\} \\ & \cup \{r_1^Z, \dots, r_m^Z\} \\ & \cup \{r_X, r_Y, r_Z, r_*\} \end{aligned}$$

and a set of users U composed of U_X , U_Y , U_Z and U^* , where, for all $\omega \in \{X, Y, Z, *\}$, we note $U_\omega = \{u_1^\omega, \dots, u_n^\omega\}$. Then the authorization policy A is informally constructed as follows: for each hyperedge $e_j = \{x_{i_1}, y_{i_2}, z_{i_3}\}$, user $u_{i_1}^X$ (resp. $u_{i_2}^Y$, $u_{i_3}^Z$) has access to r_j^X (resp. r_j^Y , r_j^Z) and to r_X (resp. r_Y , r_Z), and user u_j^* has access to all resources in P but r_j^X , r_j^Y , r_j^Z , r_X , r_Y and r_Z . More formally, we have:

$$\begin{aligned} A = & \{(u_i^X, r_j^X) : x_i \text{ belongs to } e_j, \forall i \in [n], \forall j \in [m]\} \\ & \cup \{(u_i^Y, r_j^Y) : y_i \text{ belongs to } e_j, \forall i \in [n], \forall j \in [m]\} \\ & \cup \{(u_i^Z, r_j^Z) : z_i \text{ belongs to } e_j, \forall i \in [n], \forall j \in [m]\} \\ & \cup \{(u_j^*, r_h^\omega) : \forall \omega \in \{X, Y, Z\}, \forall j, h \in [m], j \neq h\} \\ & \cup \{(u_i^\omega, r_\omega) : \forall i \in [n], \forall \omega \in \{X, Y, Z\}\} \\ & \cup \{(u_j^*, r_*) : \forall j \in [m]\} \end{aligned}$$

To conclude the construction, which can be done in polynomial time, we set $d = k$, and the resiliency policy is thus $\text{res}(P, 0, d, 4)$.

First, suppose that there exists a solution M' for the 3-DIMENSIONAL MATCHING problem. Without loss of generality, assume that $|M'| = k$, $M' = \{e_1, \dots, e_k\}$, and that $e_i = (x_i, y_i, z_i)$ for all $i \in [k]$ (recall that all members of M' are pairwise disjoint). Then, observe that for all $i \in [k]$, user u_i^* has access to all resources but r_i^X , r_i^Y , r_i^Z , r_X , r_Y and r_Z . However, u_i^X has access to r_i^X and r_X , user u_i^Y has access to r_i^Y and r_Y , and user u_i^Z has access to r_i^Z and r_Z . Hence, we have $N(\{u_i^X, u_i^Y, u_i^Z, u_i^*\}) = P$, and, since all members of M' are pairwise disjoint, we thus constructed a set of teams for $\text{RCP}\langle s = 0, t = 4 \rangle$, or, in other words, $\text{res}(P, 0, d, 4)$ is satisfiable.

Conversely, suppose that there exist V_1, \dots, V_d , pairwise disjoint subsets of U such that for all $i \in [d]$, we have $|V_i| = 4$ and $N(V_i) = P$. We first claim that for all $i \in [d]$, V_i intersects U_X (resp. U_Y , U_Z and U_*) on exactly one element. Indeed, otherwise, since $|V_i| = 4$ and since all users in U_X (resp. U_Y , U_Z , U_*)

have access to only r_X (resp. r_Y, r_Z, r_*) among $\{r_X, r_Y, r_Z, r_*\}$, V_i could not have access to all these resources. Thus, we know that for all $i \in [d]$, we have $V_i = \{u_{i_1}^X, u_{i_2}^Y, u_{i_3}^Z, u_{i_4}^*\}$, for some $(i_1, i_2, i_3, i_4) \in [n] \times [n] \times [n] \times [m]$. We claim that $(x_{i_1}, y_{i_2}, z_{i_3}) = e_{i_4}$. Indeed, observe that user $u_{i_4}^*$ has access to all resources but $r_{i_4}^X, r_{i_4}^Y, r_{i_4}^Z, r_X, r_Y$ and r_Z . By construction, the only way for having $N(V_i) = P$ is that user $u_{i_1}^X$ (resp. $u_{i_2}^Y, u_{i_3}^Z$) has access to resources $r_{i_4}^X$ (resp. $r_{i_4}^Y, r_{i_4}^Z$) or, in other words, that x_{i_1} (resp. y_{i_2}, z_{i_3}) belongs to hyperedge e_{i_4} . Thus, there exists k pairwise disjoint hyperedges in M . \square

3 Refined positive results for the case $s = 0$

We now turn to the particular case where $s = 0$. As said in Section 1, one motivation for studying this case is that it is the bottleneck of the algorithm of Li *et al.* [10] for $\text{RCP}\langle \rangle$. Hence, we believe that designing efficient algorithms for this sub-case might help us solve much larger instances of $\text{RCP}\langle \rangle$ than is currently possible. To this end, we now provide a complete characterization of the complexity when considering all possible combinations of parameters among p, d and t . We also investigate the question of reduction rules within the framework of kernelization, highlighting a difference of behavior between $\text{RCP}\langle s = 0 \rangle$ and $\text{RCP}\langle s = 0, t = \infty \rangle$.

3.1 FPT algorithms

The first algorithm is a dynamic programming-based approach similar to the one for SET COVER [4], in order to obtain an FPT algorithm for $\text{RCP}\langle s = 0 \rangle$ parameterized by (p, d) . While this result was already known, given that $\text{RCP}\langle \rangle$ is itself FPT with this parameterization (and that $\text{RCP}\langle s = 0 \rangle$ is actually FPT parameterized by p only, as we will see in Theorem 8), we provide for $\text{RCP}\langle s = 0 \rangle$ a better running time. In particular, as we will see later, a previous known reduction of Li *et al.* [10] actually proves that when d is fixed, the obtained running time is the best we can hope for, under the Exponential Time Hypothesis (ETH)³.

Theorem 6. $\text{RCP}\langle s = 0 \rangle$ can be solved in $O^*(2^{dp})$ time.

Proof (of Theorem 6). Let $U = \{u_1, \dots, u_n\}$. We define a dynamic programming algorithm which, given any $i \in [n]$ and any d -tuple of subsets of P (S_1, \dots, S_d) , returns *yes* if there exist d mutually disjoint sets T_1, \dots, T_d , each being a subset of $\{u_1, \dots, u_i\}$ and such that $S_j \subseteq N(T_j)$ for all $j \in [d]$, and returns *no* otherwise (in which case we will say that such an algorithm is *correct*). To do so, we define the following recursive formula DP . First, we set:

$$\begin{aligned} DP(0, S_1, \dots, S_d) &= 1 \text{ if and only if } S_j = \emptyset, \text{ for all } j \in [d]; \\ DP(i, \emptyset, \dots, \emptyset) &= 1 \text{ for all } i \in [n]. \end{aligned}$$

³ The ETH claims that SAT cannot be solved in $O^*(2^{o(n)})$, where n is the number of variables in the CNF formula [7].

For the induction, let $i \in [n]$ and $\mathcal{S} = (S_1, \dots, S_d)$ where $S_j \subseteq P$ for all $j \in [d]$. Let $J = \{j \in [d] : S_j \neq \emptyset\}$, and for all $j \in J$, define $\mathcal{S}_j = (S_1, \dots, S_{j-1}, S_j \setminus N(u_i), S_{j+1}, \dots, S_d)$. Finally, we set:

$$DP(i, \mathcal{S}) = DP(i-1, \mathcal{S}) \vee \left(\bigvee_{j \in J} DP(i-1, \mathcal{S}_j) \right)$$

Lemma 1. $DP(i, \mathcal{S})$ is correct.

Proof. Suppose that T_1, \dots, T_d are d mutually disjoint subsets of $\{u_1, \dots, u_i\}$ such that $S_j \subseteq N(T_j)$ for all $j \in [d]$. We may assume that $T_j \neq \emptyset$ iff $S_j \neq \emptyset$. Then, either $u_i \notin T_j$ for all $j \in [d]$, in which case $DP(i-1, \mathcal{S})$ returns *yes*, or $u_i \in T_j$ for some $j \in [d]$, which implies $S_j \neq \emptyset$ and thus $j \in J$. In this case $DP(i-1, \mathcal{S}_j)$ returns *yes*.

Conversely, if $DP(i-1, \mathcal{S})$ return *yes*, then there exist d mutually disjoint sets T_1, \dots, T_d , each being a subset of $\{u_1, \dots, u_{i-1}\}$ (and thus a subset of $\{u_1, \dots, u_i\}$), and such that $S_j \subseteq N(T_j)$ for all $j \in [d]$. If $DP(i-1, \mathcal{S}_j)$ returns *yes* for some $j \in J$, then there exist d mutually disjoint sets T_1, \dots, T_d , each being a subset of $\{u_1, \dots, u_i\}$ and such that $S_q \subseteq N(T_q)$ for all $q \in [d]$, $q \neq j$, and $S_j \setminus N(u_i) \subseteq N(T_j)$. In this case $S_j \subseteq N(T_j \cup \{u_i\})$. \square

Clearly, $DP(n, P, \dots, P)$ returns *yes* if and only if $\text{res}(P, 0, d, t)$ is satisfiable. A table of size $n2^{dp}$ is sufficient to store all intermediate results, while each step takes $O(d)$ time, establishing the claimed running time. \square

Li *et al.* [10] showed that $\text{RCP}\langle s = 0, t = \infty, d = 3 \rangle$ is NP-hard, by a reduction from 3-DOMATIC PARTITION, which transforms a graph of n vertices into an instance $(U, R, UR, \text{res}(P, 0, 3, \infty))$ with $|P| = n$. Since a $2^{o(n)}$ algorithm for 3-DOMATIC PARTITION would violate the ETH (by a linear reduction from SAT [1]), we have the following:

Theorem 7. $\text{RCP}\langle s = 0, t = \infty, d = 3 \rangle$ cannot be solved in $2^{o(p)}$ time unless the ETH fails.

Hence, for fixed d , the algorithm described in Theorem 6 has an optimal running time. We continue our quest for a better understanding of the frontier between tractable and intractable cases of the $\text{RCP}\langle s = 0 \rangle$ problem. Given the positive result parameterized by (p, d) , a natural question is to consider each parameter separately. The question can well be answered negatively concerning the parameter d , since, as we saw before, $\text{RCP}\langle s = 0, d = 3, t = \infty \rangle$ is NP-hard [10], and thus $\text{RCP}\langle s = 0 \rangle$ is para-NP-hard parameterized by d . However, we are able to give a different answer for the parameter p only.

Theorem 8. $\text{RCP}\langle s = 0 \rangle$ is FPT when parameterized by p .

Proof. The result makes use of Lenstra's celebrated algorithm [9] for Integer Linear Programming Feasibility (ILPF) parameterized by the number of variables.

Theorem 9 (Lenstra [9]). *Whether a given ILP has a non-empty solution set can be decided in $O^*(f(n))$ time for some computable function f , where n denotes the number of variables of the ILP.*

Note that this algorithm has been improved by Kannan [8], with $f(n) = n^{O(n)}$ (but exponential space), and by Frank and Tardos [5] so that the algorithm runs in polynomial space, and with $f(n) = O(n^{2.5n+o(n)})$.

We thus give an ILPF formulation of the problem with a number of variables depending on p and t . As we saw previously, since we may assume that $t \leq p$ in any positive instance, the result will follow (by Lenstra's result) for the parameterization by p only.

Let $(U, R, UR, \text{res}(P, 0, d, t))$ be the input instance of $\text{RCP}(s = 0)$. For any $N \subseteq P$, let U_N denote the set of users having neighborhood exactly N in P , or, formally: $U_N = \{u \in U : N(u) = N\}$. Moreover, we define the following set called *configurations*:

$$\mathcal{C} = \left\{ \{N_1, \dots, N_b\} : b \leq t, N_i \subseteq P, i \in [b], \bigcup_{i=1}^b N_i = P \right\}.$$

For any $N \subseteq P$, we note

$$\mathcal{C}_N = \{c = \{N_1, \dots, N_{b_c}\} \in \mathcal{C} : N = N_i \text{ for some } i \in [b_c]\}$$

the set of configurations involving N . Informally, a configuration $\{N_1, \dots, N_b\}$ represents a way to dominate P , by picking one user in U_{N_i} , for each $i \in [b]$.

The variables of our ILP are in one-to-one correspondence with elements of \mathcal{C} , and will be denoted by $\{x_c : c \in \mathcal{C}\}$. Since \mathcal{C} is of size bounded by $O(\sum_{b=1}^t 2^{bp})$, the number of variables is bounded by a function of p and t only. Then, we define the following two sets of constraints:

1. $\sum_{c \in \mathcal{C}} x_c = d$,
2. $\sum_{c \in \mathcal{C}_N} x_c \leq |U_N|$ for all $N \subseteq P$.

We now explain the idea of the ILP. Observe that in a positive instance, there always exists a set of teams in which in each set, each user has a different neighborhood. For any $T \subseteq U$, define $\phi(T) = \{N(u) : u \in T\}$, the set of neighborhoods of users in T . Then, by definition of the problem, for any set of teams $V = \{T_1, \dots, T_d\}$, we have $\phi(T_i) \in \mathcal{C}$ for all $i \in [d]$. Notice that we might have $\phi(T_i) = \phi(T_j)$ for $i, j \in [d]$, $i \neq j$. We can associate, with each such set of teams, a vector $X^V = \{x_c^V\}_{c \in \mathcal{C}}$, where x_c^V is the number of sets of V having configuration $c \in \mathcal{C}$. By the remark above, we might have $X^V = X^{V'}$ for two different sets of teams V and V' , in which case we will say that these two sets of teams are *configuration-equivalent*. Observe that given a vector $X = \{x_c\}_{c \in \mathcal{C}}$ such that $X = X^{V^*}$ for a fixed set of teams V^* , we can construct in polynomial time a set of teams V that is configuration-equivalent to V^* ; constraints (1) and (2) aim to find such a vector. Suppose that there exists a set of teams $V^* = \{T_1, \dots, T_d\}$ of the problem. It is clear that X^{V^*} fulfills constraints (1) and (2). Conversely, constraints in (1) ensure that the set of teams will contain d

sets, while constraints in (2) ensure that when constructing a set of configuration $c = \{N_1, \dots, N_{b_c}\}$, there must exist a new user having neighborhood exactly N_i for all $i \in [b_c]$ and that has not been already assigned to another set. \square

3.2 User reductions

We now focus on reduction rules which can be performed in polynomial time and result in an equivalent instance having a smaller number of users. More formally, we say that a (decision) problem has a *kernel* [4] of size f , for some computable function $f : \mathbb{N} \rightarrow \mathbb{N}$, if there exists a polynomial algorithm which, given an instance x with parameter k , outputs an instance x' of size $|x'|$ with parameter k' such that: (i) $k' \leq k$, (ii) x is positive if and only if x' is positive, and (iii) $|x'| \leq f(k)$. In the case of $\text{RCP}\langle s = 0 \rangle$ our aim is thus to obtain an equivalent instance with a number of users bounded by a function of d and t .

While the role of t was so far of less interest for the complexity of the problem, we show that the problem behaves differently from the kernelization point of view, depending on whether $t = \infty$ or not. We first show that when $t = \infty$, the problem admits a kernel with at most dp users. To do so, we will make use of the following:

Lemma 2 (*d*-expansion Lemma [2]). *Let $d \geq 1$ be a positive integer and $G = (A, B, E)$ be a bipartite graph with bipartition (A, B) and $E \subseteq A \times B$ such that for all $b \in B$, $N(b) \neq \emptyset$. If $|B| \geq d|A|$, then there exist non-empty vertex sets $X \subseteq A$ and $Y \subseteq B$ which can be found in time polynomial in the size of G , such that:*

- (i) $N(Y) \subseteq X$, and
- (ii) *there is a d -expansion of X into Y : a collection $M \subseteq E \cap (X \times Y)$ such that every vertex of X is incident to exactly d edges of M , and exactly $d|X|$ vertices of Y are incident to an edge of M .*

Theorem 10. $\text{RCP}\langle s = 0, t = \infty \rangle$ admits a kernel with at most dp users.

Proof. Suppose we are given an instance of $\text{RCP}\langle s = 0, t = \infty \rangle$. We present two reduction rules which are used to decrease the number of users. For each of these rules, we will prove that the instance is positive iff the reduced instance is positive, in which case we will say that the rule is *safe*.

Reduction Rule 1: if there exists $u \in U$ with $N(u) = \emptyset$, then delete u .

Proof (of safeness). Simply observe that such a user cannot participate in any set of teams if the instance is positive, and, conversely, cannot turn a negative instance into a positive one if it is deleted. \square

Reduction Rule 2: if there exist $X \subseteq P$, $Y \subseteq U$ such that $N(Y) \subseteq X$ and there is a d -expansion of X into Y , then delete X from P , Y from U , and $(Y \times X) \cap UR$ from UR .

Proof (of safeness). If the instance is a positive one, then there exists a set of teams $\{V_1, \dots, V_d\}$. Then, for all $r \in P \setminus X$, there does not exist $u \in Y$ such that $(u, r) \in UR$, since $N(Y) \subseteq X$. Hence, $N(V_i \setminus Y) \supseteq P \setminus X$, and thus $\{V_1 \setminus Y, \dots, V_d \setminus Y\}$ is a set of teams for the reduced instance, which is thus a positive one.

Conversely, suppose that the reduced instance is a positive one: there exist V_1, \dots, V_d , disjoint sets of users from $U \setminus Y$ such that $N(V_i) \supseteq P \setminus X$. Since there is a d -expansion of X into Y , for all $r \in X$, there exist $u_1^r, \dots, u_d^r \in Y$ such that $(u_i^r, r) \in UR$ for all $i \in [d]$, where $u_i^r \neq u_{i'}^{r'}$ for all $r \neq r'$ and $i \neq i'$. Hence, for all $i \in [d]$, if we set $V_i' = V_i \cup \{u_i^r : r \in X\}$, we have $V_i' \cap V_j' = \emptyset$ for all $1 \leq i < j \leq d$, and $N(V_i') \supseteq P$ for all $i \in [d]$, and thus we have a positive instance as well, which proves that the rule is safe. \square

Since each reduction rule can be applied in polynomial time, and since each of them decreases the number of users by at least one, the algorithm runs in polynomial time. Finally, by Lemma 2, if none of the previous reduction rules applies, then $|U| \leq dp$, and we thus have a kernel with at most dp users, as desired. \square

As Li *et al.* [10] point out, $\text{RCP}\langle s = 0, d = 1 \rangle$ is equivalent to the SET COVER PROBLEM. Known kernel lower bounds for this problem [3] lead to the following theorem, which is in sharp contrast to the previous case.

Theorem 11. $\text{RCP}\langle s = 0, d = 1 \rangle$ (and thus $\text{RCP}\langle s = 0 \rangle$) does not admit a kernel with $(p + t)^{O(1)}$ users, unless $\text{coNP} \subseteq \text{NP}/\text{poly}$.

4 Conclusion and future work

We considered $\text{RCP}\langle \rangle$, a problem introduced recently in the area of access control to analyze the resiliency of a system. Given the large number of natural parameters in an instance of this problem, and given that these parameters are likely to take small values in practice, our goal was to provide a systematic analysis of the complexity of the problem using the framework of parameterized complexity. For all but one possible combination of the parameters, we were able to obtain either a positive or negative result. We also considered a restricted variant of the problem for which we settled the parameterized complexity of all possible combinations of the parameters. A first obvious idea of future work is thus to fill the remaining hole of Figure 1, namely to decide whether $\text{RCP}\langle \rangle$ is in FPT, XP, W[1]-hard or para-(co)NP-hard parameterized by p .

Another interesting further line of research would be to study resiliency aspects with respect to other problems. In the context of graphs for instance, we could define the problem of determining whether upon removal of at most s vertices, a given graph still satisfies some property given by another combinatorial problem, *e.g.* having a vertex cover of size k . We believe that considering structural parameterizations (together with s) might lead to interesting new results. As in our case, the complexity of such a new problem will certainly depend on the complexity of the considered underlying problem (*i.e.* the case $s = 0$).

References

1. Nadia Creignou. The class of problems that are linearly equivalent to satisfiability or a uniform method for proving NP-completeness. *Theoretical Computer Science*, 145(1-2):111 – 145, 1995.
2. Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
3. Michael Dom, Daniel Lokshtanov, and Saket Saurabh. Incompressibility through colors and IDs. In *proceedings of ICALP'09*, pages 378–389, 2009.
4. Rod G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
5. András Frank and Éva Tardos. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica*, 7(1):49–65, 1987.
6. Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
7. Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
8. Ravi Kannan. Minkowski’s convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, August 1987.
9. Hendrik W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.
10. Ninghui Li and Mahesh V. Tripunitara Qihua Wang. Resiliency policies in access control. *ACM Trans. Inf. Syst. Secur.*, 12(4), 2009.
11. Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, 1996.
12. Qihua Wang and Ninghui Li. Satisfiability and resiliency in workflow authorization systems. *ACM Trans. Inf. Syst. Secur.*, 13(4):40, 2010.